

# Reference Manual





# **REFERENCE MANUAL for TMAS Common-View Time and Frequency Measurement System**

---

## **TABLE OF CONTENTS**

<b>I.</b>	<b>Introduction</b>	<b>2</b>
<b>II.</b>	<b>Installation</b>	<b>3</b>
<b>III.</b>	<b>Configuring the System</b>	<b>5</b>
<b>IV.</b>	<b>Antenna Position</b>	<b>8</b>
<b>V.</b>	<b>Making Measurements</b>	<b>10</b>
<b>VI.</b>	<b>Viewing Data on the Web</b>	<b>17</b>
<b>VII.</b>	<b>NIST Disciplined Oscillator</b>	<b>18</b>

The calibration was performed by making a common-clock comparison to the reference system at NIST for a period of 5 days using a different antenna and antenna cable than the one to be used when installed at the customer's site. This calibration determined that the receiver delay was approximately 40 ns. The customer needs to enter an Rx delay value that is equivalent to the propagation delay through their antenna cable plus 40 ns.

---

### **Certichron (Weehawken) NIST-TMAS System Information**

<b>System ID:</b>	63
<b>Result of Rx delay calibration:</b>	40 ns (add propagation delay of antenna cable to this number)
<b>Web Link for comparisons:</b>	<a href="http://gps.nist.gov/control/certichron.htm">http://gps.nist.gov/control/certichron.htm</a>
<b>User Name:</b>	cvgps063
<b>Password:</b>	SIM063

## **I. Introduction**

The TMAS common-view time and frequency measurement system is designed to continually measure a laboratory's primary time and frequency standard with respect to UTC(NIST), the United States national standard for time and frequency. All comparisons are made using the multi-channel common-view GPS measurement technique. Each participating laboratory is required to connect a stable 5 or 10 MHz signal to the system to serve as a time base, and to provide a 1 Hz signal from their primary standard or UTC time scale. If the NIST Disciplined Oscillator option (Section VII) is installed, both the time base signal and the 1 Hz signal are provided by the supplied rubidium oscillator.

The system does not analyze or graph the data; it simply collects and stores it. The collected data is sent to a Web server located at NIST using the Internet file transfer protocol (FTP). The client can view the raw data and time difference graphs using a standard web browser.

The following instructions explain how to use the system. The instructions are divided into three sections: installation, operation, and viewing data on the Web.

This system should perform as described here, but if not, please report all problems to NIST so that updates to the software and/or documentation can be made if necessary.

## II. Installation

The TMAS computer requires six input connections:

- **GPS Antenna (TNC cable)**

Connect the GPS antenna cable to the TNC connector on the rear panel of the system that is labeled “GPS ANTENNA”. If the customer is using their own GPS antenna and cable it must supply a 1575.42 MHz signal to the receiver. If the signal is downconverted, it must be upconverted back to 1575 MHz before being connected to the receiver.

- **Reference signal (1 Hz)**

Connect a 1 Hz signal from the supplied rubidium oscillator to the BNC connector on the rear panel labeled “REFERENCE (1 PPS)”.

- **External Time Base (BNC, must be 5 or 10 MHz)**

Connect a 10 MHz signal from the supplied rubidium oscillator (or from another frequency standard) to the BNC connector labeled “TIMEBASE” on the rear panel. The signal level should be between 200 mV and 3.5 V peak-to-peak, or 70 mV to 1.25 V (rms) with 50  $\Omega$  termination.

- **Serial Interface cable**

Connect one end of the supplied DB-9 serial null modem cable to the rubidium oscillator. Connect the other end to the top DB-9 connector on the rear panel of the computer.

- **KVM Switch**

Connect the system to the Startech KVM switch and provide NIST with log-in information.

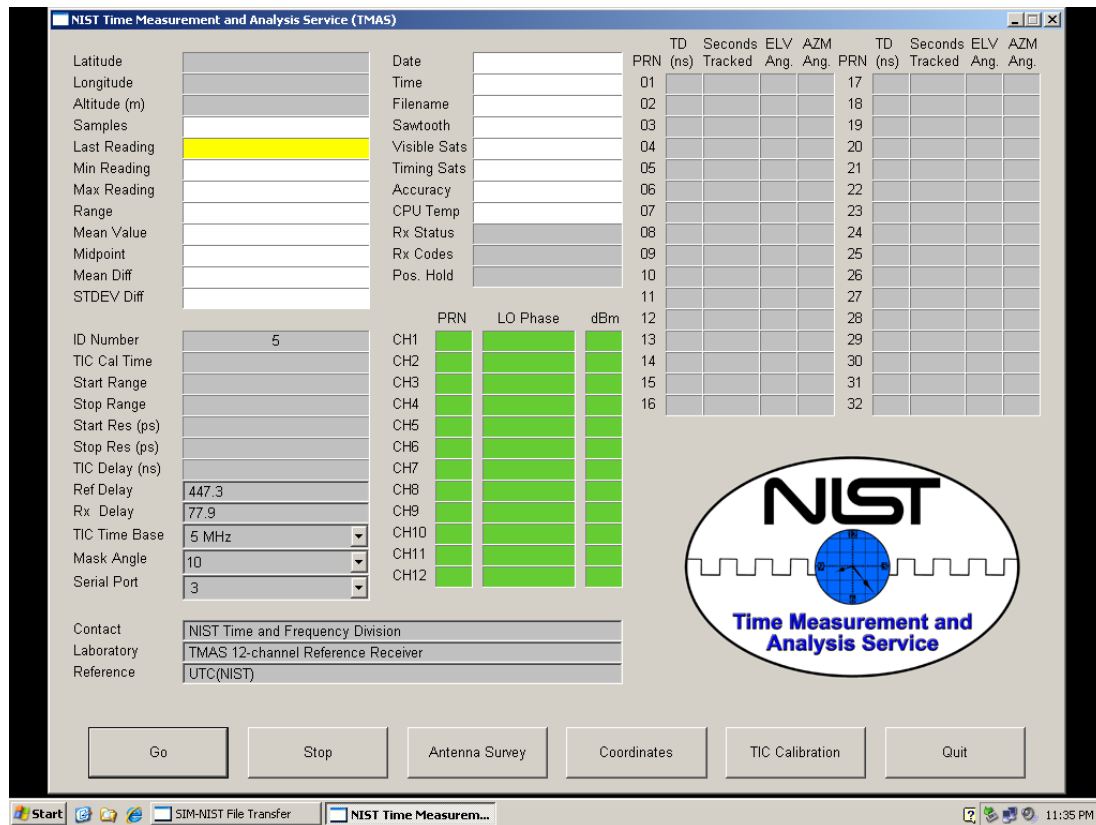
- **Network Connection**

The system contains an Ethernet card for connection to the Internet. Provide your dedicated IP address to NIST, and perform the following steps:

- Click on the START button, then click Settings, then Network Connection
- Doubleclick on Local Area Connection, then click on Properties
- Doubleclick on Internet Protocol (TCP/IP)
- Enter the IP address, Subnet mask, Default gateway and DNS servers

### III. Configuring the System

After the computer is turned on and Windows has finished loading, the TMAS measurement software should run automatically. You will then see a display similar to Figure III.1 below:



*Figure III.1. Initial Screen Display*

You will also see another application on the Windows taskbar called SIM\_NIST\_FTP. This is the file server software that sends the current data file back to a web server located at NIST every 10 minutes. Normally, you should just ignore this application and let it run. If you wish, you can click on it to see the name of the current data file and to see the time and result of the last FTP attempt (an attempt is made every 10 minutes). You can even click the middle button to force the current file to transfer right now. However, do not click the STOP button, because this software has to run at all times. If you accidentally click the STOP button, run the SIM\_NIST\_FTP program from the \SIM directory again. **Remember, unless SIM\_NIST\_FTP is running, no data will be uploaded to the network so it is critical that this application remains running at all times.**

A third application will appear on the toolbar if the NIST Disciplined Oscillator option (Section VII) has been installed.

Some of the fields displayed near the bottom left corner of the TMAS measurement display might already contain information, but this information needs to be verified and changed if necessary before making measurements. Note that the “ID Number” cannot be changed without editing a system file. It should not be changed under any circumstances, because it is the unique number that identifies your system.

To make changes, point to and click on the field that you want to edit. The fields are described below:

Ref Delay	This number should be entered in units of nanoseconds, with 0.1 ns resolution. It is the estimated delay in the cable that connects the laboratory’s primary standard (1 Hz signal) to the measurement system. If a distribution amplifier is used, it must include the delay in the distribution system. Cable delays can usually be measured with an uncertainty of about 1 nanosecond using a time interval counter. In the case of the NISTDO, it will be a short cable with a small delay. If the cable is supplied by NIST, the delay will already be entered.
Rx Delay	This number is usually entered at NIST in units of nanoseconds and should not be changed. It represents the combined delay of the GPS receiver, antenna, and antenna cable. NIST supplies the Rx delay number based on the result of the calibration, and usually ships the measurement system with the same antenna and cable used during the calibration. If the customer is using their own antenna and cable, see page 1 for information about the number to enter.
TIC Time Base	This field allows you to select the frequency of your time base oscillator (either 5 MHz or 10 MHz). The system will not work unless this setting matches the frequency of your oscillator
Mask Angle	The mask angle determines the elevation angle above the horizon where satellites will be tracked. It can be set from 0 to 25° in 5° increments. For most receiving locations, a mask angle of 10° is recommended. A mask angle lower than 10° can reduce the quality of the timing solution, and raising the mask angle reduces the number of visible satellites.
Serial Port	This selects the number of the COM port used to control the GPS receiver. It has been preset at NIST and should not need to be changed.

Contact	The name of the group or division within the laboratory or organization.
Laboratory	The name of the laboratory or organization.
Reference	The manufacturer and model number of the primary standard, for example, “NISTDO”, or the name of the laboratory’s primary time scale, for example, “UTC(NIST)”.



## IV. Antenna Position

After the configuration information has been entered, the next step is to establish the GPS antenna coordinates. This can be done either by entering the coordinates if they are previously known, or by allowing the system to survey its own antenna position. Once the antenna position is known it is saved by the system and used for all further measurements, so this process does not need to be repeated unless the antenna is moved. The uncertainty of the antenna position directly contributes to the uncertainty of the time and frequency measurements recorded by the system, so it is important to use the best antenna coordinates that you have available. Note that if the system surveys its own position, the latitude and longitude values are usually very accurate (within 1 meter), but the altitude error can sometimes exceed 10 meters, so an independent survey of altitude normally reduces the uncertainty.

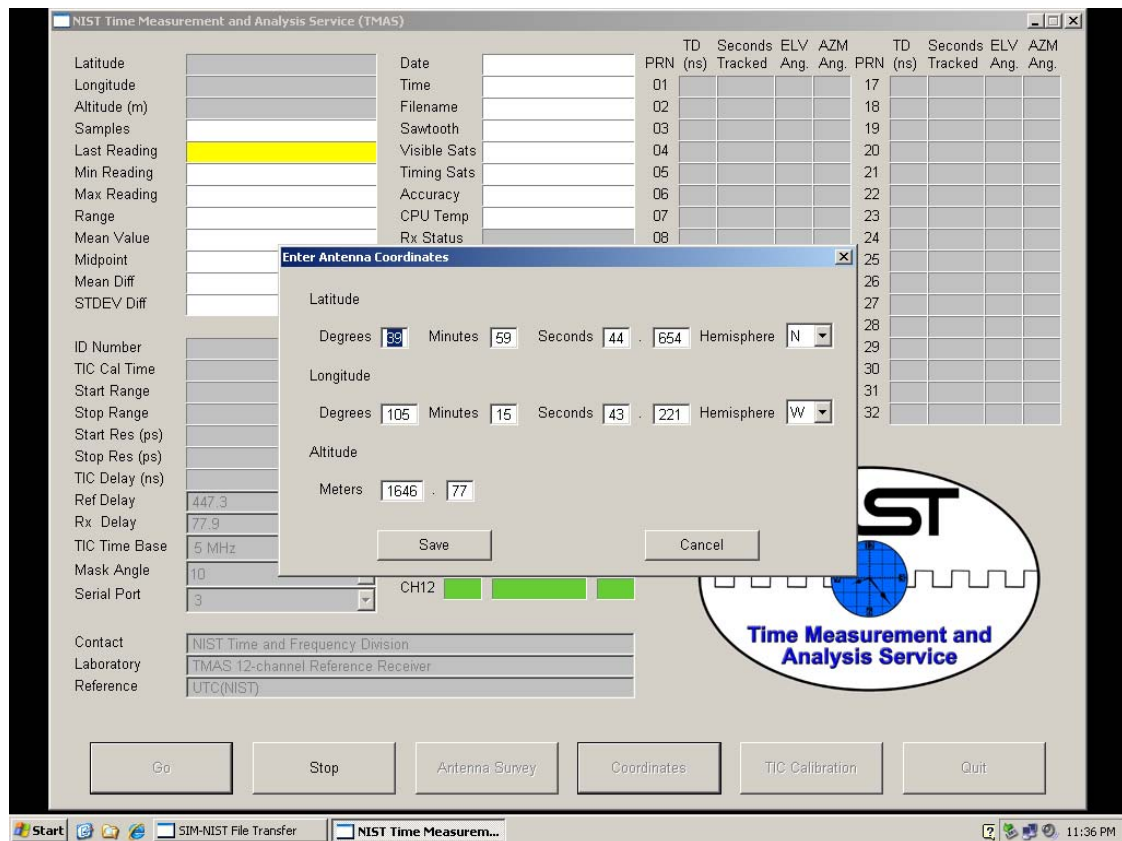


Figure IV.1. The Enter Coordinates Window

To manually enter the antenna position:

- Click the *Coordinates* button. A window appears as shown in Figure III.1.

- Enter the antenna coordinates. The allowable resolution is 1 milliarcsecond for latitude and longitude, and 1 centimeter for altitude.
- When finished, click the *Save* button.

To have the system survey its antenna position:

- Click the *Antenna Survey* button.
- The receiver will be reset and begin to look for satellites. It might take up to 20 minutes before it produces its first position fix. Once the first position fix is obtained, coordinates are averaged for 24 hours (86,400 seconds), so the entire antenna survey takes slightly more than 24 hours to complete.
- During the antenna survey, the latitude, longitude, and altitude fields will be updated, and the samples field will show the number of position fixes that have been averaged so far.
- After 86,400 valid position fixes have been obtained, the average antenna position is saved and the system is ready to begin measurements.

## V. Making Measurements

After the antenna position data is available, click the *Go* button to start the measurements. At this point, the system will calibrate the time interval counter, check the status of the GPS receiver, and then begin to collect and store measurements. When the measurements have begun, the screen display will look similar to Figure V.1. Leave the system running in this mode with the measurement screen displayed at all times. If it is ever necessary to stop the measurements, click the *Stop* button.

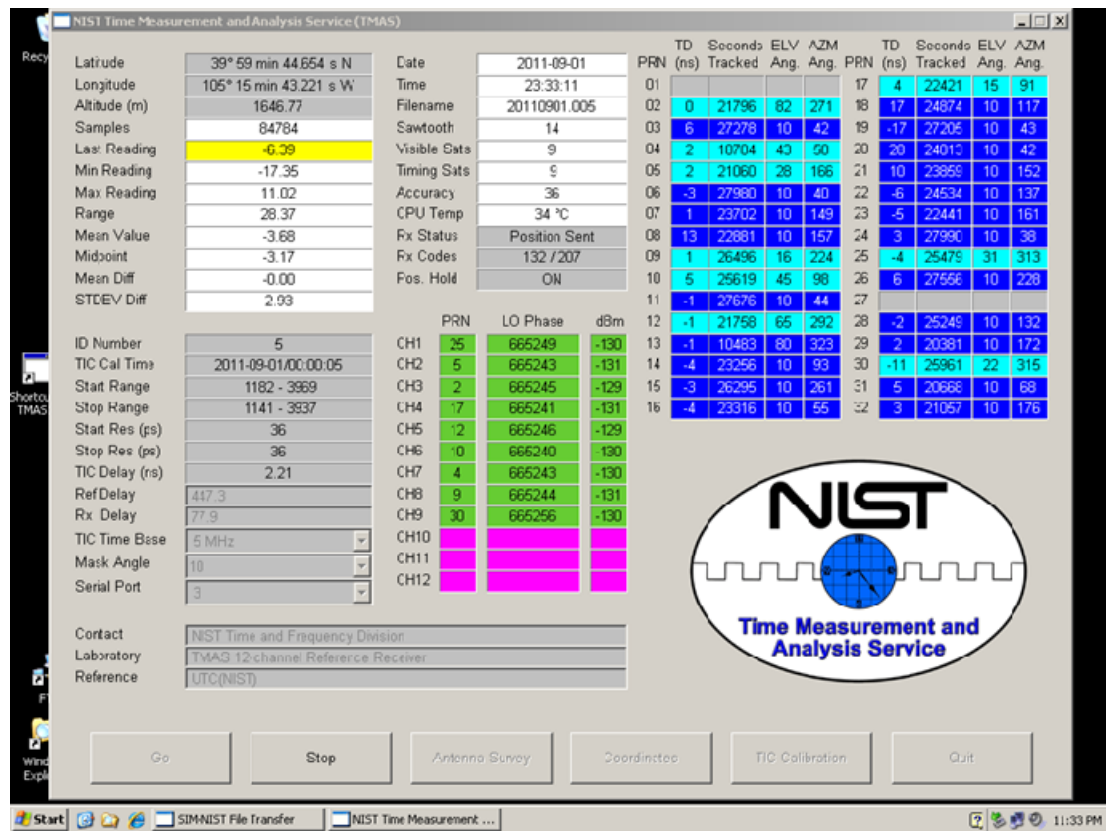


Figure V.1. The Measurement Screen

The measurement screen is divided into three areas (left, middle, and right). The following tables describe what the numbers in each area mean.

*Table V.1. Left Area of Measurement Screen*

<b>Table Heading</b>	<b>Description of Item</b>
Latitude	The latitude of the GPS antenna, as determined by the antenna survey or as entered by the user. The available resolution is 1 milliarcsecond.
Longitude	The longitude of the GPS antenna, as determined by the antenna survey or as entered by the user. The available resolution is 1 milliarcsecond.
Altitude	The altitude of the GPS antenna, as determined by the antenna survey or as entered by the user. The available resolution is 1 centimeter.
Samples	The number of data samples recorded so far during the current day. One sample is recorded every second, so 86400 samples are recorded per day.
Last Reading	The most recent reading from the time interval counter (displayed on a yellow background). The unit is nanoseconds. If the timing pulse from the GPS receiver is early with respect to the local standard, a negative number will be displayed. This is normal.
Min Reading	The smallest reading recorded from the time interval counter during the current day. The unit is nanoseconds.
Max Reading	The largest reading recording from the time interval counter during the current day. The unit is nanoseconds.
Range	The maximum reading minus the minimum reading, in nanoseconds. This gives an indication of the worst case timing uncertainty of the system during the current day.
Mean Value	The total of all readings divided by the number of samples, in nanoseconds.
Midpoint	The midpoint of all readings, in nanoseconds.
Mean Diff	The average difference between two successive readings, given in nanoseconds. This number will often be close to zero.
STDEV Diff	The standard deviation of the difference between two successive readings, given in nanoseconds. This number provides a rough estimate of the timing stability of the system at an averaging time of 1 second.

TIC Cal Time	The date and time when the time interval counter (TIC) was last calibrated. The TIC is calibrated immediately prior to a measurement run and at the beginning of each UTC day. During the calibration, this field will display the messages “IN PROGRESS” and “TRANSITION”. If the calibration fails, that will also be indicated in this field as “CALIBRATION FAILED”.
Start Range	The counting range (min to max) of the TIC’s start interpolator during a 100 nanosecond interval, displayed only for diagnostic purposes.
Stop Range	The counting range (min to max) of the TIC’s stop interpolator during a 100 nanosecond interval, displayed only for diagnostic purposes.
Start Res (ps)	The resolution of the start input on the time interval counter, given in picoseconds. It is equal to the start range (max – min) divided by 100 nanoseconds. The resolution should be 30 picoseconds or less, and should be nearly equivalent to “Stop Res”.
Stop Res (ps)	The resolution of the stop input on the time interval counter, given in picoseconds. It is equal to the stop range (max – min) divided by 100 nanoseconds. The resolution should be 30 picoseconds or less, and should be nearly equivalent to “Start Res”.
TIC Delay (ns)	The time offset due to delays in the time interval counter, given in nanoseconds. This number is used as a correction value that is applied to each reading. It is usually less than 0.5 ns.

*Table V.2. Middle Area of Measurement Screen*

<b>Table Heading</b>	<b>Description of Item</b>
Date	The current date read from the GPS receiver.
Time	The current time read from the GPS receiver. Due to the screen updating procedure, the displayed time has an uncertainty of $\pm 1$ second.
Filename	The filename currently being created by the system. The format for the name is YYYYMMDD. The file extension is the system ID number. One file is created per day.
Sawtooth	The timing pulse produced by the receiver is ambiguous to local oscillator clock cycles, due to the nature of the receiver hardware. However, the receiver firmware keeps track of the offset of the next timing pulse, and that number is displayed here and applied by the system as a “sawtooth” correction. The correction ranges from $-13$ to $+13$ ns.
Visible Sats	The number of satellites currently visible to the GPS receiver. In rare instances, this number can be higher than 12, but the receiver is only capable of tracking 12 satellites at one time.
Timing Sats	The number of satellites currently being tracked and used in the timing solution. The maximum number is 12.
CPU Temp	The temperature obtained from a sensor on the CPU board inside the computer case, near the GPS receiver. This sensor is read every second, but the resolution is $1\text{ }^{\circ}\text{C}$ , so the readings generally do not change rapidly. The temperature tends to be a few degrees higher than the laboratory temperature, and to fluctuate as the laboratory temperature fluctuates. For best results, it is important to not let the temperature change rapidly by keeping the receiver in a room where the temperature does not fluctuate by more than $\pm 3\text{ }^{\circ}\text{C}$ .
Rx Status	This field contains a text message that normally only changes during signal acquisition or an antenna survey. When the system is taking readings, the displayed message will read “Position Sent”.
Rx Code	This field contains a number ranging from 1 to 255 that indicates the receiver status. If the receiver is operating

	properly, the number should always be an 8 or a 32. Other numbers will appear during the acquisition or antenna survey process and can be ignored.
Pos. Hold	ON means that the system is using a fixed antenna position, and is not trying to calculate further position fixes. OFF will only appear during signal acquisition or an antenna survey.
Box with green background	<p>The “green” box contains information about each of the 8 channels on the GPS receiver. If one the channels is not being used, the box is blank with a light red background.</p> <p>The “PRN” column contains the psuedo random noise code (PRN) for each satellite being tracked. PRN codes have possible values ranging from 1 to 32.</p> <p>The “LO Phase” column shows the current phase (in nanoseconds) of the GPS receiver’s local oscillator (LO) with respect to the received GPS signal. These numbers change rapidly, by more than 10 microseconds per second, because the LO has an intentionally introduced frequency offset that exceeds <math>1 \times 10^{-5}</math>. This offset is removed when the timing solution is produced. The numbers are displayed only for diagnostic purposes and to show that the receiver is tracking satellites.</p> <p>The “dBm” column shows the signal strength of each satellite being tracked. These numbers should normally be near -130 dBm. Numbers smaller than -135 dBm on all channels could indicate that the antenna cable is too long or that local signal conditions are poor.</p>

The right area of the screen provides data collected from all GPS satellites in view. The data is organized in row numbers labeled with the PRN numbers from the satellites. The 32 rows represent all possible slots in the GPS constellation, and it is not likely that all possible slots will contain satellites.

*Table V.3. Right Area of Measurement Screen*

<b>Table Heading</b>	<b>Description of Item</b>
TD (ns)	<p>The time difference (in nanoseconds) between the last reading recorded from the specified satellite and the average value of all satellites in view.</p> <p>If the background on this field is light blue, it means that the satellite is currently part of the timing solution, and this field is updating every second. If the background of this field is dark blue, it is not currently updating, and it contains the last recorded time estimate from the satellite. In some cases, the last recorded time estimate might be a fairly large number, which is why the satellite was dropped from the timing solution.</p>
Seconds Tracked	<p>The number of samples collected from the specified PRN during the current UTC day. The GPS satellites have an orbital period of approximately 12 hours, and therefore pass over a given location on earth twice each day. The receiver tracks each satellite down to the selected mask angle. With a 10° mask angle it is not uncommon to collect more than 400 minutes (24000 seconds) of data from a satellite each day. Not all satellites will be tracked on both of their passes.</p> <p>If the background on this field is light blue, it means that the satellite is currently part of the timing solution, and this field is updating every second. If the background of this field is dark blue, it is not currently updating, but this indicates the satellite was being tracked earlier in the UTC day.</p>
ELV Ang.	<p>Normally, the values shown on a dark blue background will be equivalent to the mask angle, since satellites are dropped from the timing solution when their elevation angle is lower than the mask angle. However, in cases where more than 8 visible satellites are available for selection, or when a specific satellite is producing bad data, the receiver might drop a satellite from the solution when it is still well above the mask angle. Therefore, you will sometimes see elevation angles high than the mask angle printed on a dark blue background.</p>
AZM Ang.	<p>If the background of this field is light blue, it shows the current</p>



azimuth angle of the satellite being tracked. If the background of the field is dark blue, it shows the azimuth angle of the satellite at the moment when it was dropped from the timing solution.

**NOTE:** Once the system is operational, it will continue to take measurements unless there is a power outage, or if it is stopped by the operator. The measurement screen will reset at the end of each UTC day, erasing the old information and beginning the collection of new information. If the system is stopped for any reason, you can restart it by clicking on the TMAS icon on the desktop, and then clicking the *Go* button.

## **VI. Viewing the Data on the Web**

When the system is running properly, it will send data every 10 minutes to a web server located at NIST. This data can then be viewed using any web browser with new data made available every 10 minutes. The web software allows up to 200 days of data to be displayed at one time.

To view the collected data, visit this link:

<http://gps.nist.gov/control/certichron.htm>

User name: cvgps063

Password: SIM063

The page allows you to view common-view GPS comparisons between your time standard and UTC(NIST) in Boulder, and to view one-way GPS data recorded at either site.

## VII. NIST Disciplined Oscillator

The NIST Disciplined Oscillator (NISTDO) option is included with this TMAS unit. The NISTDO software disciplines a Symmetricom 8040C rubidium oscillator to agree with the UTC(NIST) time scale in Boulder, Colorado. A 9-pin serial null modem cable is used to connect the rubidium oscillator to the TMAS unit.


The NISTDO software will run automatically when the TMAS unit is turned on. To start the NISTDO, simply click on the GO button.

The NISTDO works by utilizing both the Internet and “common-view” observations of Global Positioning System (GPS) satellites, and can serve as the primary frequency and time standard for a calibration or metrology laboratory. It is directly referenced to the Coordinated Universal Time scale kept at NIST, known as UTC(NIST). This makes it easy to establish traceability to the International System (SI) directly through NIST. Customers are provided with standard frequency outputs of 5 MHz and/or 10 MHz, as well as 1 pulse per second timing outputs. The NISTDO outputs are accurate to within about  $\pm 20$  ns (peak-to-peak variation) with respect to UTC(NIST) and provide frequency with an uncertainty near  $5 \times 10^{-14}$  when averaged over a 24-hour interval. The NISTDO runs automatically, but the following sections describe the theory of operation for those who are interested.

### VII.1. Theory of Operation

The NISTDO is based on the CVGPS technique, a simple but effective method for comparing two clocks. Ideally, a comparison between two clocks would be made by bringing them both to the same location and making a direct comparison. However, when the two clocks are located at different sites, the time difference between them can still be measured by simultaneously comparing both clocks to a signal that is in “common-view” of both sites. The difference between the two comparisons is the time difference between the two clocks. The common-view signal is simply a vehicle used to transfer time from one site to the other. Its accuracy is unimportant because it does not influence the final measurement result.

The CVGPS method involves a GPS satellite ( $S$ ), and two receiving sites ( $A$  and  $B$ ), each containing a GPS receiver and a local clock (Figure VII.1). The satellite transmits a signal that is received at both  $A$  and  $B$ , and  $A$  and  $B$  each compare the received signal to their local clock. Thus, the measurement at site  $A$  compares the GPS signal received over the path  $d_{SA}$  to the local clock,  $Clock A - S$ . Site  $B$  receives GPS over the path  $d_{SB}$  and measures  $Clock B - S$ .



*Figure VII.1. The common-view GPS measurement method.*

The difference between the two measurements is an estimate of *Clock A – Clock B*. Delays that are common to both paths  $d_{SA}$  and  $d_{SB}$  cancel even if they are unknown, but uncorrected delay differences between the two paths add uncertainty to the measurement result. Thus, the basic equation for a CVGPS measurement is

$$(Clock\ A - S) - (Clock\ B - S) = (Clock\ A - Clock\ B) + (e_{SA} - e_{SB}).$$

The components that make up the  $(e_{SA} - e_{SB})$  error term include delay differences between the two sites caused by ionospheric and tropospheric delays, multipath signal reflections, environmental conditions, or errors in the GPS antenna coordinates. These factors can be measured or estimated and applied as a correction to the measurement, or they can be accounted for in the uncertainty analysis. It is also necessary to calibrate the GPS receivers used at both sites and account for the delays in the receiver, antenna, and antenna cable.

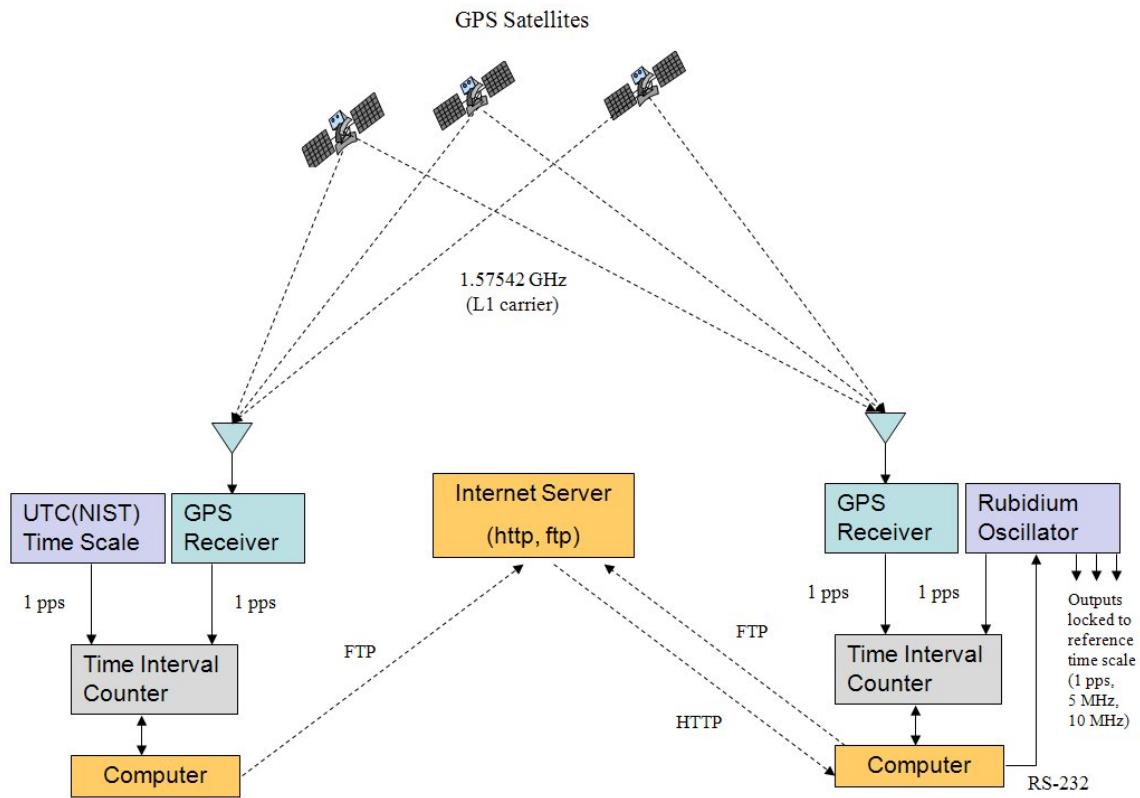


Figure VII.2. The NISTDO system.

Figure VII.2 is a diagram of the NISTDO system. One common-view GPS system is located at NIST, the second at the customer's facility. These systems are supplied by NIST to its TMAS customers. Each system includes an eight-channel GPS receiver (C/A code, L1-band) and a time interval counter. The NIST system compares a 1 pulse per second (pps) timing signal from the GPS receiver to the UTC(NIST) time scale. The

customer's system compares a 1 pps signal from the GPS receiver to the rubidium oscillator supplied by NIST.

The measurement systems at both sites average time interval counter readings for 10 minutes and then simultaneously upload their results to an Internet file transfer protocol (FTP) server. The use of FTP requires transmission control protocol (TCP) ports 20 and 21 to be left open on the local firewalls. After the data are uploaded, the NISTDO invokes a common gateway interface (CGI) applet on the Internet server that instantly processes the CVGPS data. This applet, called *CVDIFF*, aligns and differences data from the individual satellite tracks, and discards data collected from satellites that are not in common view at both sites. The average time difference, *TD*, between the clocks at the two sites is obtained by

$$TD = \frac{\sum_{i=1}^N (REFGPS_i(A) - REFGPS_i(B))}{N}, \quad (2)$$

where *N* is the number of satellites tracked by both GPS receivers, *REFGPS<sub>i</sub>(A)* is the series of individual satellite tracks recorded at the customer's site, and *REFGPS<sub>i</sub>(B)* is the series of tracks recorded at NIST.

The server includes another applet, called *AVDIFF*, for use by customers that are located so far away from NIST that few if any satellites are in common-view. *AVDIFF* implements the "all-in-view" method, where the satellite tracks are not aligned and no tracks are discarded. Instead, the averages of the *REFGPS<sub>i</sub>(A)* and *REFGPS<sub>i</sub>(B)* data series are calculated, and the time difference *TD* is the difference between the two averages.

Both *CVDIFF* and *AVDIFF* send data through TCP port 80, where it can be read by the NISTDO by use of the hypertext transfer protocol (HTTP). Thus, the NISTDO can nearly instantly obtain the time difference between its local oscillator and the reference time scale, and apply this information to discipline the local rubidium oscillator.

A PID controller was chosen to discipline the rubidium. Its purpose is simply to correct the error, *e*, between a measured process variable and a desired set point (*SP*). Here the process variable is *TD*, the last measured time difference between the local oscillator and UTC(NIST). Because the NISTDO is attempting to lock the local oscillator as closely as possible to UTC(NIST), the desired value of *SP* is 0.

A new value for *TD* is obtained every 10 minutes and a steering correction is sent to the rubidium oscillator. The steering correction is always a dimensionless frequency correction, and time errors are corrected through frequency adjustments. The NISTDO software (Figure VII.3) displays the latest values of *TD* in a scrolling window, and also displays other information related to the PID controller. The tuning parameters for the PID controller can be changed, but this should not be done without first consulting NIST.

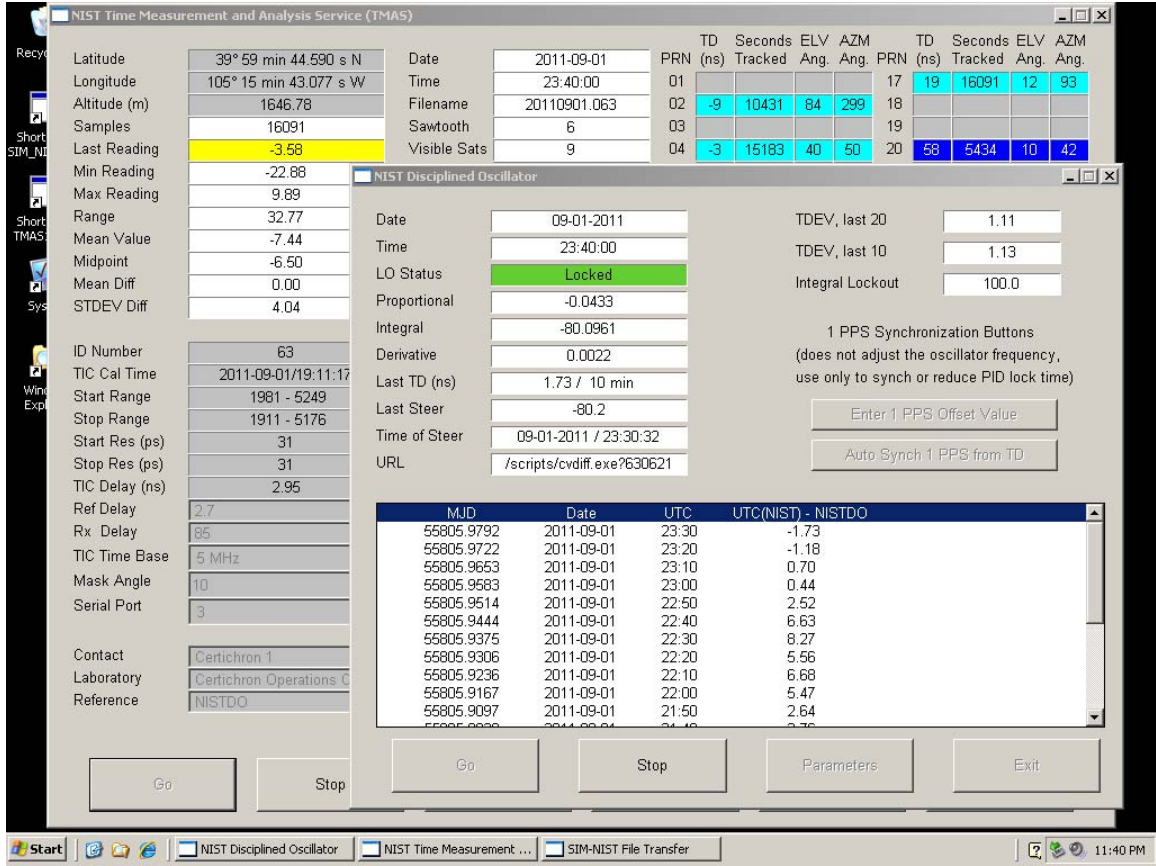


Figure VII.3. NISTDO software running on a TMAS system.

The NISTDO software also displays whether the CVDO is “locked” or “unlocked.” The CVDO is considered to be locked if its output is both accurate and stable with respect to the reference. Two criteria must be met to satisfy the lock condition: the most recent time difference must be less than 50 ns (accuracy) and the time deviation,  $\sigma_x(\tau)$ , of a series of the recent time differences must be less than 10 ns at  $\tau = 10$  minutes (stability). The time deviation is a metric for time stability based on the modified Allan deviation,  $Mod \sigma_y(\tau)$ , and is computed as


$$\sigma_x(\tau) = \frac{\tau}{\sqrt{3}} Mod \sigma_y(\tau) \quad .$$

### VII.3. NISTDO Performance

The NISTDO is configured to have three 10 MHz outputs, one 5 MHz output, and two 1 pulse per second (pps) outputs. However, one of the 1 pps outputs is dedicated to the common-view measurements.

The NISTDO is steered every 10 minutes. A 10-minute update period provides the quickest response to an unlocked condition and minimizes the deviation from the set point. For most applications, this advantage outweighs the slight increase in phase noise caused by the additional steering.

Figure VII.4 is a phase plot of a NISTDO compared to UTC(NIST) for the 75-day period ending on April 25, 2010. The data points are 1-hour averages. Note that the average time offset of the NISTDO with respect to UTC(NIST) was near zero (0.07 ns) with only a few outliers falling more than 15 ns from the mean. The phase plot has essentially no slope or trend and thus the frequency offset is negligible, less than  $1 \times 10^{-17}$ .



*Figure VII.4. NISTDO phase compared to UTC(NIST).*

Figure VII.5 shows the frequency and time stability,  $Mod \sigma_y(\tau)$  and  $\sigma_x(\tau)$  respectively, of the NISTDO's 1 pps output with respect to UTC(NIST) at intervals of 10 minutes and longer. The frequency stability reaches  $1 \times 10^{-12}$  after less than two hours of averaging and drops to  $6 \times 10^{-15}$  at  $\tau = 1$  day. The time stability is near or below 1 ns after a few hours of averaging. After about 10 days of averaging, the frequency stability is near  $1 \times 10^{-16}$ .






Figure VII.5. Frequency stability and time stability (diagonal lines) of NISTDO with respect to UTC(NIST).

### VII.3. NISTDO Failure Modes

Several situations can cause a NISTDO to fail or become unlocked. Like a GPS disciplined oscillator (GPSDO), a NISTDO is vulnerable to GPS outages due to local interference or other causes. The problem is more pronounced with a NISTDO, however, because a GPS failure either at NIST or at the customer's site can cause a failure. In addition, a NISTDO is vulnerable to Internet outages at either NIST or the customer's site.

If an Internet or GPS outage is long enough it will eventually cause the NISTDO to fail. However, short outages are normally not a problem. The rubidium oscillator is tuned very close to its nominal frequency while locked, and will continue to keep accurate time without steering corrections for a reasonably long interval. Internet and/or GPS outages of up to about one hour should not be noticeable, and time can be kept within a few microseconds of UTC(NIST) for one day or longer even if both the Internet and GPS are unavailable.

### VII.4. Summary

A NIST disciplined oscillator (NISTDO) is a unique new instrument that makes it possible for calibration and metrology laboratories to maintain a standard that is both synchronized and syntonized to UTC(NIST), the national standard for time and frequency in the United States.